THE PRICING OF SOVEREIGN RISK UNDER COSTLY INFORMATION*

September 15th, 2017

Grace Weishi Gu† and Zachary R. Stangebye‡

We explore in a quantitative model the consequences of costly information acquisition for the pricing of sovereign risk. We develop an approach to identify information costs empirically using Google search data. The calibrated model delivers a number of interesting results: First, it endogenously generates levels of time-varying volatility in the country risk spread that approach the data; second, it suggests that risk premia fluctuate with country-specific states, which has econometric consequences; and third, it highlights a non-trivial welfare trade-off in the promotion of fiscal transparency.

Keywords: costly information, sovereign default, time-varying volatility, risk premia, transparency

JEL code: F34, D83, G12

*The authors would like to thank Manuel Amador, Mark Aguiar, Sebnem Kalemli-Ozcan, Boragan Aruoba, Luminita Stevens, Felipe Saffie, John Shea, Juan Carlos Conesa, Marina Azzimonti, Ting Liu, Raoul Minetti, Charles Engel, Kamran Bilir, Kenneth West, Leonardo Martinez, Sandra Lizarazo, Ruediger Bachmann, Olivier Jeanne, Michal Szkup, Mark Wright, Marco Bassetto, Linda Tesar, Pablo Ottonello, Gita Gopinath and seminar participants at the University of Maryland, Michigan State University, Stony Brook University, the University of Pittsburgh, the University of Wisconsin, the International Monetary Fund, the Federal Reserve Bank of Chicago, the University of Michigan, and the NBER Summer Institute for useful comments and suggestions. All errors and mistakes are our own.

†E-mail: grace.gu@ucsc.edu, Phone: (831) 459-4791, Address: 463 Engineering 2, 1156 High Street, Santa Cruz, CA 95064

‡E-mail: zstangeb@nd.edu, Phone: (574) 631-1054, Fax: (574) 631-1054, Address: 3015 Nanovic Hall, Notre Dame, IN 46556 (Corresponding Author)

1. Introduction

Yields in sovereign bond markets in emerging economies largely reflect the risk that a domestic borrower may default on foreign creditors since it does not in principle have any reason to care about the well-being of these creditors. But a sovereign borrower's lack of welfare concern for its foreign lenders is not the only relevant friction that arises from the international nature of these markets: Information frictions also play a large role in cross-border financial transactions: Investors are likely to be less informed about payoff-relevant shocks in other countries (Hatchondo [2004], Van Nieuwerburgh and Veldkamp [2009], or Bacchetta and van Wincoop [2010]). Information about such shocks can often be acquired, but typically at a cost.[1]

In this paper we seek to understand how costly information acquisition affects the equilibrium pricing of sovereign default risk. In particular, we construct a model in which the sovereign's default and borrowing decisions as well as the lenders' acquisition of payoff-relevant information are jointly endogenous: Lenders always observe some public states, such as output growth and debt levels, but cannot directly observe other potentially payoff-relevant states when making investment decisions, such as the severity of a recession implied by a potential future default or populist sentiment in that country. Information regarding these sorts of shocks can only be acquired at a cost.

For simplicity of exposition, we will colloquially say that 'investors/lenders acquire' information, though our market is a bit more nuanced. In any market-

[1] Cole and Kehoe (1998), Sandleris (2008), Catao et al. (2009), and Pouzo and Presno (2015) have all shown that information asymmetries are key to explaining various features of these markets, though none have considered the consequences of allowing information to be gathered at a cost.

based model of costly information acquisition, one encounters the dilemma highlighted by Grossman (1976), which is that market prices tend to reveal too much information and thus kill incentives to acquire information in the first place. We circumnavigate this issue by separating the information acquisition decision from market participation, much like Angeletos and Werning (2006). An independent contractor interested only in the integrity of its forecast such as the credit rating agencies in Holden et al. (2017) or Manso (2013), conducts the information acquisition. As market participants, lenders have costless access to the forecasts provided by the contractor and pay attention to them when they convey useful information. Real world analogues of the forecaster might be financial software and analytics firms, such as Bloomberg or Reuters; alternatively they could be public, supernational entities such as the IMF or the World Bank, or credit rating agencies such as Moody's or S&P.

Costly information acquisition generates a dependence of investor attention to unobserved states on those that are publicly observed. For instance, when debt levels are low and output growth is high, there is little to gain in terms of accurately inferring unobserved shocks since default risk is negligible for most of their realizations. However, for moderately high debt levels and low growth, information is more valuable since unobserved shocks may substantially affect default risk.

Intuitively, investors will start poring over more information during crises to carefully study the borrower and its default risk, e.g., professional forecasts, IMF staff reports, credit rating agency reports, and public finance records. Such practical techniques are employed in the financial sector, where fund managers in charge of multiple portfolios pay relatively limited attention to individual country risk until publicly available indicators in that country trigger some preset alert.[2]

To quantify the impact of this mechanism, we develop a novel strategy to empirically identify the magnitude of information costs using Google's Search Volume Index (SVI). The literature has proposed many different measures of information acquisition (Barber and Odean [2008], Gervais et al. [2001], and Seasholes and Wu [2007]), but SVI is one of the few *direct* measures of investor "attention." Da et al. (2011) demonstrate that it is an effective measure of attention to firm valuation and stock prices, but to the authors' knowledge the index has not yet been used to measure attention to a sovereign nation's financial position.[3]

In this paper, we match the fraction of quarters in which intense attention is paid to the borrower country, using SVI as an empirical proxy for attention in the data. If information is infinitely costly, this fraction will be zero; if it is free, this fraction will be one. The attention fraction uniformly increases as information cost decreases; thus, the latter cleanly identifies the former. We calibrate our benchmark model to Ukraine from 2004-2014, though we could perform our quantitative exercise for any country.

The model generates three new sets of results. First, it serves as a microfoundation for time-varying volatility in the country risk spread, since it generates this feature endogenously without assuming that any fundamental processes exhibit it. During normal times, lenders receive little to no information regarding payoff-relevant unobserved shocks. Consequently, they consider

[2] e.g. https://www.parametricportfolio.com/solutions/equities/emerging-markets

[3] In Appendix D, we test the robustness of the Google search index against other potential attention metrics. Investor attention by alternative metrics is strongly correlated with the Google data.

these shocks to be at their mean for the purpose of inferring and thus pricing default risk. This implies that bond yields do not respond to realizations of these unobserved shocks during normal times and so spread volatility is lower. As crises approach, however, lenders acquire additional information about these shocks, which then become priced. This implies that bond yields *do respond* to realizations of unobserved shocks during crisis times, which increases spread volatility.

This time-variation in the macroeconomic volatility is a well-documented empirical fact (Justiniano and Primiceri [2008] or Bloom [2009]) but little has been done as of yet to understand its causes,[4] despite the fact that Fernández-Villaverde et al. (2011) show that second-moment fluctuations in the risk-spread can have substantial first-moment effects on investment and output. To quantify this channel we develop a model-free metric of time-varying volatility, which we call the *Crisis Volatility Ratio* (CVR). We find that in the data the CVR is 3.67. Our calibrated benchmark model, without targeting this metric, can explain roughly 78% of it by generating a CVR of 2.86. And it does so without assuming any time-variation in the volatility of any underlying shocks. This is a substantial improvement over, for example a standard sovereign default model in the absence of costly information, which generates a CVR of 1.27.

Second, the state-contingent nature of information acquisition translates into bond risk premia. The spread on a short-term sovereign bond can be

[4]Some notable recent exceptions are Seoane (2015) and Johri et al. (2015), but these papers explain time-varying volatility in the country risk spread by assuming exogenous time-varying volatility in fundamentals. Sedlacek (2016) provides an alternative solution, but it generates time-varying cross-sectional dispersion among firms. Chari and Kehoe (2003) also show that investor herding behavior can increase uncertainty during crises.

decomposed into two components: Default risk and a risk premium for that default risk. When information acquisition is costly, the relative contribution of each to the overall spread will change across publicly observed states, e.g., debt levels and output growth. During non-crisis times, lenders lack information regarding unobserved states; the risk premium reflects this ignorance. When a crisis occurs, however, so too does the accuracy of payoff-relevant information, which works to mitigate risk exposure. Thus, while the level of the default risk is greater, lenders' effective risk-aversion is also lower. This implies that bond risk premia, independently of the level of default risk, depend on country-specific observables.

The calibrated model suggests that during crisis times default risk constitutes about 23.7% of crisis-level spreads[5], whereas a standard model without costly information acquisition would put this figure substantially lower, at about 13.6%. During non-crisis times, the average difference across these models is a mere 0.73 percentage points, so the effect is strongly state-contingent. This has substantial implications for the inference of default risk from spread data, which is a common practice in the literature (Bi and Traum [2012], Bocola [2016], Bocola and Dovis [2016], and Stangebye [2015]). In particular, it implies that a standard sovereign default model may *understate* default risk during crises by assuming a risk premium that is too large.

Third, we find a non-monotonicity in sovereign welfare as a function of information costs. The intuition is simple and operates through debt prices. If there is no transparency i.e. large information costs, then lenders demand

[5] We define a 'crisis' directly from spreads following Aguiar et al. (2016a).

a greater risk premium, making it more expensive to borrow and service debt. This risk premium falls as transparency increases, which is consistent with findings in the empirical literature (Kopits and Craig [1998], Poterba and Rueben [1999], Bernoth and Wolff [2008], and Iara and Wolff [2014]). This is tantamount to a reduction in borrowing costs for the sovereign, which improves its welfare.

However, when the sovereign becomes too transparent a risk-shifting begins to occur: At high debt levels, the sovereign will be fully exposed to the volatility that results from the pricing of normally ignored unobserved shocks. The calibrated model suggests that sovereign welfare as a function of information costs features a U-shape. Initially, increasing transparency hurts the sovereign since price volatility rises. Eventually, though, the welfare benefits of reduced borrowing costs begin to dominate and welfare attains its maximum level in the full-information model.

It is important to note that we are focusing on information frictions that are *country-specific* i.e. those that arise between a single country and its lenders. This has substantial implications for our results. For instance, one cannot account for all the time-varying volatility or the state-contingent composition of the risk-spread by controlling for global metrics such as the CBOE VIX or the P/E ratio, as has been done in the literature (Bocola and Dovis [2016] and Aguiar et al. [2016a]). With regard to the time-varying volatility, our finding corroborates the careful empirical work of Fernández-Villaverde et al. (2011), who find that the bulk of the time-varying interest rate volatility in emerging markets is country-specific rather than global.

Our focus on relations between a single borrower and its lenders over time and the implications for the ergodic price distribution distinguishes our analysis from the related work of Cole et al. (2016). These authors also explore a model of costly information acquisition in sovereign debt markets, but their focus is static. They highlight the potential for this channel to cause contagion effects across many countries and generate multiplicity. Angeletos and Werning (2006) and Carlson and Hale (2006) also explore how market-based information acquisition or rating agencies affect equilibrium multiplicity or uniqueness in variations on the canonical model of Morris and Shin (1998). Durdu et al. (2013) explore the impact of news shocks in a similar model. Recently, Bassetto and Galli (2017) also study the role of information on bond pricing in a two-period Bayesian trading game, focusing on implications for inflation risk.

The remainder of this paper is divided as follows: Section 2 describes the model; Section 3 discusses the data, quantitative implementation of the model, counterfactual analysis, as well as the model's novel implications; and Section 4 concludes.

2. Model

We consider a small open economy model of endogenous sovereign default in the vein of Eaton and Gersovitz (1981). This is in part for tractability and in part to demonstrate our model's applicability to the recent, expanding quantitative literature, e.g., Aguiar and Gopinath (2006), Arellano (2008), Hatchondo and Martinez (2009), or Mendoza and Yue (2012).[6] There is a

[6]For simplicity of exposition, as the above papers we focus on Markov Perfect Equilibria that can be expressed recursively, though the set of equilibria is potentially much larger (Passadore and Xandri [2015]).

sovereign borrower who issues one-period non-state-contingent debt to a unit mass of foreign lenders. This borrower lacks the ability to commit to repay this debt in subsequent periods and will default if it is optimal to do so ex post.

For clarity, we distinguish a random variable from its realization by placing a tilde over the former.

2.1. Shocks

There are two shocks in this model, both to the sovereign. The first is a growth shock to output that the sovereign receives each period. More specifically, the endowment Y_t can be expressed in terms of a sequence of growth rates g_s as follows:

$$Y_t = Y_0 \times \Pi_{s=1}^{t} e^{g_s}$$

where Y_0 is given. We assume g_t follows an AR(1) process:

$$g_t = (1-\rho)\mu_g + \rho g_{t-1} + \sigma_g \epsilon_t$$

where μ_g is average growth rate, ϵ_t is a standard normal, and σ_g is the standard deviation of the growth innovation. The endowment and its growth processes are publicly observed by everyone. We collect the publicly observed exogenous states into a vector $s_t = \{Y_t, g_t\}$. This vector follows a Markov process with transition density $f_t(s_{t+1}|s_t)$.

The second shock in the model is an iid default cost shock, m_t, which applies only in the first period of a default.[7] It is the only shock that is unobserved by foreign lenders when they make investment decisions (we will detail the timing below); information regarding it can only be acquired at a cost. As a default cost shock, it could represent the magnitude of private capital outflows, the impact of likely fiscal consolidation, or the severity of the ensuing international sanctions, among other default-induced costs about which lenders may not be perfectly informed. In Appendix A, we also show that when m_t is normally distributed it can alternatively be interpreted as a default preference shock, such as populist sentiment or political change. In this sense, m_t can stand in for a wide variety of payoff-relevant factors that are not immediately observed by the lenders when they make investment decisions.

2.2. Timing

The timing of events can be found in Figure 1 and is as follows: Period t begins with the realization of s_t, following which the sovereign makes a default decision. Conditional on repayment, it then chooses a level of debt issuance B_{t+1} to maximize its expected utility prior to the realization of m_{t+1}.

Next, a professional forecaster, who observes the public states s_t and B_{t+1}, chooses the accuracy with which he acquires information about m_{t+1} given some information cost. He designs a signal of the unobserved shock, x_t, and can pay a cost to increase its accuracy.

Following the information acquisition decision, m_{t+1} and x_t are jointly realized in the middle of period t. The market coordinates on the signal: Com-

[7] Allowing the shock to be applied to more than the first period of a default does not change our mechanism. This assumption is merely for tractability.

Figure 1: Timing of Events in a Period

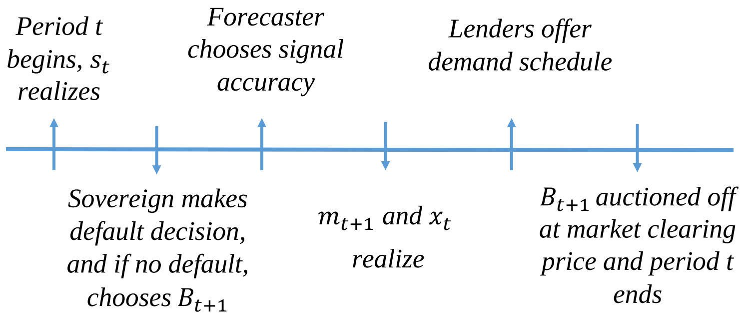

petitive lenders know both the signal and its accuracy. The sovereign then determines an issuance price that clears the bond market and period t ends.

Notice that we assume that the sovereign cannot change its bond supply following the realization of m_{t+1}.[8] This allows us to focus on the role of information acquisition and avoid the complicated and, for our purposes, unnecessary signaling game that would ensue.

2.3. Sovereign Borrower

As is standard in the literature, we use a recursive, Markov-Perfect specification with limited commitment on the part of the sovereign. At the beginning of each period, it compares the value of repaying debt, $V_{R,t}$, with that of default, $V_{D,t}$, and chooses the option that provides a greater value:

$$V_t(s_t, B_t, m_t) = \max\{V_{R,t}(s_t, B_t), V_{D,t}(s_t, m_t)\}$$

Given the timing assumption, we can express the value of repayment at the

[8]We do not specify whether sovereign gets to know m_{t+1} before or after the lenders do. Since the sovereign decides its bond supply before m_{t+1} realizes and cannot change his issuance decision in either case, it does not matter.

beginning of period t as follows:

$$V_{R,t}(s_t, B_t) = \max_{B_{t+1}\in\mathcal{B}_t} \; E_{\tilde{m}_{t+1},\tilde{x}_t}\left[\log\left(C_t(\tilde{x}_t)\right) \;\; + \beta E_{\tilde{s}_{t+1}|s_t} V_{t+1}(\tilde{s}_{t+1}, B_{t+1}, \tilde{m}_{t+1})\right]$$

$$\text{subject to} \;\; C_t(\tilde{x}_t) = Y_t - B_t + q_t(B_{t+1}|s_t, \tilde{x}_t)B_{t+1} \tag{1}$$

The sovereign has time-separable log-preferences over consumption[9] and is a monopolist in his own debt market. The determination of the issuance price schedule, $q_t(B_{t+1}|s_t, x_t)$, will be discussed in the market clearing section below. If the budget set is ever empty, either before or after the realization of x_t, we follow the literature standard and assume that the sovereign defaults.[10]

We assume that when a default happens in period t all debt is wiped out. In the first period of the default regime, the sovereign is subject to the shock m_t realized in period $t-1$. During the entire default regime, the sovereign is excluded from capital markets for a random number of periods and faces persistent output losses, though it is subject to the m_t shock only in the first period of the default.

The costs during the period of exclusion are output losses that consist of two components, known and unknown. We will denote the former with $\psi > 0$; the latter will be the shock m_t. These costs could be interpreted as the usual consequences of tightening credit conditions, a disruption of trade credit, or a banking slump (Mendoza and Yue [2012] or Sosa-Padilla [2012]). m_t is intended to stand in for the components of those costs that are harder

[9]None of the intuition behind our results relies on the assumption of log-utility. Any concave function will work. The benefit of using log-utility is that the unobserved shock can be interpreted either as an endowment/supply shock or as a preference/demand shock. We demonstrate this in Appendix A.

[10]In the calibrated model, we can easily construct the set $\mathcal{B}_t$ such that the budget set is never empty at any point in the state space and the upper bound on $\mathcal{B}_t$ never binds.

to predict, such as the magnitude of private capital outflows, the impact of likely fiscal consolidation, or the severity of the ensuing international sanctions, about which lenders may not be perfectly informed. Under these assumptions, the value of default can be expressed recursively as follows:

$$V_{D,t}(s_t, m_t) = \log\big(C_t\big) \; + \beta E_{\tilde{s}_{t+1}|s_t}\left[\phi V_{t+1}(\tilde{s}_{t+1}, 0, 0) + (1-\phi) V_{D,t+1}(\tilde{s}_{t+1}, 0)\right]$$
$$\text{subject to} \;\; C_t = Y_t \times e^{-\psi + m_t} \tag{2}$$

We assume that m_t is normally distributed around zero.[11] ϕ is the Poisson rate at which the sovereign regains access to international capital markets.

We define the sovereign's default decision with a binary operator:

$$d_t(m_t, s_t, B_t) = \mathbf{1}\{V_{R,t}(s_t, B_t) < V_{D,t}(s_t, m_t)\}$$

2.4. Forecaster

There is an inherent difficulty associated with market-based information acquisition problems: The price tends to convey too much information. Perfect Bayesian investors can infer all relevant information from market prices, which gives them no incentive to acquire information in the first place (Grossman [1976] or Dow and Gorton [2006]).

Rather than bounding rationality by including noise traders (Kyle [1985]), we circumvent this problem by designing the market to operate with complete rationality and transparency on a signal of the true hidden information

[11] The calibrated level of σ_m will be sufficiently small that 'efficient' defaults in which $m_t > \psi$ will have near zero probability and will never materialize on any simulated path.

rather than the hidden information itself.[12] Thus, no additional information regarding payoff-relevant states can be gleaned from the price besides what participants already have.

To do so, we separate the information acquisition decision from market participation, much like Angeletos and Werning (2006). We assume that all lenders are fully rational, but that each acquires information by employing the same contractor, whom we call the forecaster. In a sense, the forecaster serves as a market maker by coordinating market information. Veldkamp (2011) argues that information is a commodity that is often difficult to obtain but essentially costless to disseminate. Our market set-up reflects this feature of information production by allocating the task to a single specialist. Once the information is produced, it can be costlessly gathered by market participants.

The forecaster has a technology capable of gathering information regarding the unobserved shock, m_{t+1}, at a cost κ that we specify later. He sells his services to lenders as an independent contractor for a fixed, non-state-contingent fee, $l_t \geq 0$. We take l_t to be exogenous and assume for simplicity that lenders must pay it to have access to the sovereign bond market. After paying this fee, they have costless access to a forecaster-provided signal of the unobserved shock and will pay attention to this signal whenever they find it useful. Real-world analogues of the forecaster might be financial software and information firms such as Bloomberg or Reuters; alternatively, they could be financial news or media outlets such as the Wall Street Journal of the Financial Times; credit rating agencies such as Moody's or S&P; or multinational public institutions

[12]The other popular option is to add unseen noise to the aggregate supply e.g. Grossman and Stiglitz (1980) or Angeletos and Werning (2006). In our model, however, supply is determined by sovereign optimality. Thus, there is little plausible room for such noise.

such as the International Monetary Fund.[13]

The forecaster is interested in the integrity of its forecasts, much like the rating agencies in Holden et al. (2017) or Manso (2013). It actively weights this objective against information acquisition costs. In each period it produces a signal, x_t, of the next period's unobserved shock, m_{t+1}. The signal and the unobserved state are jointly normal, and the information contained in this signal is reflected in $\rho_{mx,t} = corr(x_t, m_{t+1}) \in [0, 1]$, which is the forecaster's choice.[14] That is, the forecaster can modify the signal to make it more or less informative about m_{t+1}: More informative signals will feature a larger $\rho_{mx,t}$. For this reason we will call $\rho_{mx,t}$ the *accuracy* or *precision* of the signal. The forecaster uses this information to publish a forecast (distribution) over all future states, observed and unobserved: $h_t(s_{t+1}, m_{t+1}|x_{t+1}, s_t) = f_t(s_{t+1}|s_t)g_{\rho_{mx,t}}(m_{t+1}|x_t)$. The goal is to minimize the mean-square-error of the default-risk forecast under this distribution.

In our benchmark, we assume m_{t+1} and x_t to be orthogonal to observed states s_t, but our framework is flexible enough to allow for some correlation with no change in the mechanism. The forecaster would simply acquire residual information that is not conveyed through observed states.

The information required to obtain a signal is given by a time-invariant function, $I(\rho_{mx,t})$, which is increasing in accuracy. The per-unit cost of information is a constant κ. In the benchmark, we assume that $I(\cdot)$ is the reduction in entropy in m_{t+1} that comes from knowledge of x_t, but our results do not

[13] Rating agencies serve a similar purpose, though they are typically in the employ of the countries whose debt they grade. Our model is isomorphic to one in which rating agencies of this type are the information gatherers so long as those fees are similarly non-state-contingent.

[14] Our restriction to signals with positive correlation is without loss of generality, since negatively correlated signals have the same information content.

hinge on this functional form.[15] Any increasing function would work.

We formulate the forecaster's information acquisition problem as below, given s_t and B_{t+1}:

$$\min_{\rho_{mx,t}\in[0,1]} \quad E_{\tilde{x}_t} E_{\tilde{m}_{t+1},\tilde{s}_{t+1}|\tilde{x}_t,s_t} \left[d_t(\tilde{m}_{t+1},\tilde{s}_{t+1},B_{t+1}) - \bar{d}_t\right]^2 + \kappa I(\rho_{mx,t}) \qquad (3)$$

$$\text{subject to} \quad \bar{d}_t = E_{\tilde{m}_{t+1},\tilde{s}_{t+1}|s_t}\left[d_t(\tilde{m}_{t+1},\tilde{s}_{t+1},B_{t+1})\right]$$

where $d_t(\cdot)$ is the binary default identifier.[16] To see the benefit of information acquisition, notice that the variance of the interior expectation is decreasing in $\rho_{mx,t}$. Consequently, the variance in the forecaster's default forecast can be reduced if he is willing to undergo costly information acquisition. When s_t and B_{t+1} indicate a greater risk of default, acquisition of more accurate information will be optimal, since the unobserved shock matters for the variance of the forecast. When these publicly known states indicate instead that there is little to no default risk, the forecaster can save on information costs and provide imprecise or even orthogonal signals (i.e., $\rho_{mx,t} = 0$), because m_{t+1} is adding little to no variance to the forecast.

2.5. Foreign Lenders

There is a unit mass of risk-averse foreign lenders who invest in risky sovereign debt. These lenders act competitively, similarly to Lizarazo (2013) or Aguiar et al. (2016a). Lenders arrive in overlapping generations and each

[15] This notion of information was developed primarily by Shannon (1958) and applied to economics by Sims (2003, 2006). Here we have $I(\rho_{mx,t}) = \frac{1}{2}\log_2\left(\frac{1}{1-\rho^2_{mx,t}}\right)$.

[16] We prefer to think of the objective as a utility rather than a resource cost, so we do not incorporate their fee, l_t, into the objective. Since the fee is non-state-contingent, doing so would not change the model.

lives for two periods. Each lender is endowed with wealth, w_t, pays the fixed contractor fee, l_t, which is assumed to be relatively small in relation to B_t or w_t, to gain access to the forecaster's signals. Then they solve a portfolio allocation problem, deciding how much to invest in risky sovereign debt and how much to invest in a risk-free asset yielding a return, r_t.

When making investment decisions in period t, lenders observe s_t and B_{t+1}. Further, they can costlessly access the forecaster's signal x_t. The lenders know the signal's informativeness since they know the policy function of the forecaster: $\rho_{mx,t} = \rho_t(s_t, B_{t+1})$. If x_t contains useful information, they will choose to include it in their information set. However, if it is uninformative i.e. $\rho_{mx,t} = 0$, out of indifference we assume that they choose not to include it in their information set.

This assumption allows us to interpret the model such that lenders' attention to signals is positively correlated with the forecaster's information acquisition. The assumption changes neither the behavior of the model nor the mechanism that lenders acquire information through the forecaster, but the interpretation will guide the model calibration; in particular, our identification of information costs.

Given the above setup, when choosing to include x_t in the information set,[17] investor i takes the bond issuance price, q, as given and solves the following

[17] For the sake of brevity we do not write out explicitly the problem without x_t in the information set.

problem:

$$\max_{b_{i,t+1}} \; E_{\tilde{s}_{t+1},\tilde{m}_{t+1}|s_t,x_t}\left[\frac{c_{i,t+1}^{1-\gamma_L}}{1-\gamma_L}\right] \tag{4}$$

$$\text{subject to} \;\; c_{i,t+1} = (w_t - l_t - b_{i,t+1}q)(1+r) + b_{i,t+1}\left[1 - d_t(\tilde{m}_{t+1}, \tilde{s}_{t+1}, B_{t+1})\right]$$

The equilibrium pricing function, as well as all other equilibrium objects, are also contained in the lenders' information set, but since it contains no hidden information there is no Bayesian extraction problem to be undertaken as in Robert E. Lucas (1972), Grossman (1976), or Bassetto and Galli (2017).

We denote aggregate bond demand in any period by $B_{D,t+1}(s_t, x_t, B_{t+1}, q) = \int_0^1 b^{\star}_{i,t+1}(s_t, x_t, B_{t+1}, q)di$. Note that there is no heterogeneity amongst lenders, and so the i-index will be irrelevant in the benchmark.

2.6. Market Clearing

Once the correlation is chosen by the forecaster, the signal is realized and distributed to all lenders, who then enter a competitive market with a common information set. The sovereign issues its predetermined debt stock, B_{t+1} at the highest possible price. Market clearing requires that $B_{D,t+1} = B_{t+1}$. This yields a pricing schedule identical to that in Aguiar et al. (2016b) but with the inclusion of the signal realization as an additional state.

2.7. Equilibrium Definition

Having described the model, we can now define our equilibrium:

Definition 1. *A* **Markov Perfect Equilibrium** *is a set of functions,* $\{V_t(s_t, B_t, m_t), V_{R,t}(s_t, B_t), A_t(s_t, B_t), V_{D,t}(s_t, m_t), q_t(B_{t+1}|s_t, x_t), \rho_t(s_t, B_{t+1})\}_{t=0}^{\infty}$

such that

1. $V_{R,t}(s_t, B_t)$ *and* $V_{D,t}(s_t, m_t)$ *solve Recursions 1 and 2 and imply the policy* $B_{t+1} = A_t(s_t, B_t)$. *Further,* $V_t(s_t, B_t, m_t) = \max\{V_{R,t}(s_t, B_t), V_{D,t}(s_t, m_t)\}$.
2. $\rho_t(s_t, B_{t+1})$ *solves Problem 3.*
3. $q_t(B_{t+1}|s_t, x_t)$ *ensures that* $B_{t+1} = B_{D,t+1}(s_t, x_t, B_{t+1}, q_t(B_{t+1}|s_t, x_t))$ *where bond demand is derived from Problem 4.*

In Appendix B, we demonstrate how this model can be stationarized for solution purposes. This model will behave in most respects as a standard quantitative sovereign default model in the vein of Eaton and Gersovitz (1981). The novel feature is how information is collected and transmitted and what that implies for bond price.

The model endogenously generates acquisition of information that is contingent on observable states. During non-crisis times, information regarding unobserved shocks is not particularly valuable for improving the default-risk forecast. Thus, such information is not acquired and sovereign debt is priced assuming these shocks to be at their average.

During crises, however, such information becomes valuable. Consequently, it is both acquired and priced, which increases the volatility of the sovereign spreads. This generates a number of interesting results including time-varying spread volatility, state-contingent risk premia that are relevant for econometric inference, and novel welfare results on transparency for the sovereign borrower, which we demonstrate below.

3. Quantitative Analysis

To determine the impact that costly information acquisition has on the pricing of sovereign risk, we calibrate the model to match a set of empirical moments from Ukraine from 2004-2014. We choose Ukraine since the volatility of its real growth rate process during this time is similar to Argentina during the 1990's, which is the canonical calibration choice for models in this vein (Aguiar and Gopinath [2006] or Arellano [2008]). Aguiar et al. (2016a, 2016b) show that growth volatility is particularly important for generating realistic spread dynamics.

Further, Ukraine was at the heart of several news cycles over the course of this period, including political upheavals during the Russo-Georgian War in August 2008 and the annexation of Crimea by the Russian Federation in early 2014. Last, we choose the period starting from 2004 because it is the only period for which Google search volume data, which is key in our approach, is available. We solve the model using value function iteration on a discrete grid. Details regarding the solution method can be found in Appendix B.

3.1. Data and Calibration

We take data from three primary sources: First is the JP Morgan Emerging Market Bond Index (EMBI) database taken from Datastream; second is the World Bank database; and third is Google Trends' Search Volume Index (SVI).

3.1.1. Information Cost Identification

First and most importantly, to find a proper cost value per unit information, i.e., κ, we match the variability of information acquisition in the model and in the data. The model has a direct measure of information acquisition, ρ_{mx}. In the data, we follow Da et al. (2011) and proxy for lenders' attention behavior with Google search trends for terms for which lenders are likely to search.

In particular, we use Google Trends' *Search Volume Index* (SVI), which is calculated as the number of searches on a particular term divided by the total searches in the time range it represents (a month in our case) and then normalized between 0 and 100. Zero implies minimum intensity and 100 implies maximum intensity. According to Google, the measure is intended to represent search interest over time; it does not convey absolute search volume.

By no means do we consider our Google search metric to be an exhaustive description of the information acquisition efforts of either the forecaster or the lenders. Rather, much like Da et al. (2011), we interpret it as a useful proxy for broader search and information acquisition activity. In Appendix D, we show that our metric is highly correlated with other indicators for lenders' or forecasters' attention, such as sovereign bond extreme returns and sovereign credit rating changes from S&P, Moody's, and Fitch.

The process of our information cost identification is as follows. First, we obtain monthly SVI data for 2004-2016 and average them into quarterly data. Then, we transform the raw SVI series into the *Abnormal Search Volume Index* series (ASVI), as suggested by Da et al. (2011). The ASVI is meant to capture a notion of paying "extra" attention to a certain event or item in a period t. In any period t, it is computed as follows:

$$ASVI_t = \log(SVI_t) - \log\left(\text{Median}\left(\{SVI_{t-9}, SVI_{t-8}, \ldots, SVI_{t-1}\}\right)\right)$$

In the benchmark, we compute the ASVI for the search term "Ukraine IMF," thus assuming the International Monetary Fund to serve the role of forecaster. We also assume $l_t = 0$, since the IMF does not receive funds directly from lenders.[18] In addition, we consider other search terms for robustness in Appendix D. The full ASVI series can be found in Figure 2. We can see that ASVI lines up well with Ukrainian crises. It also slightly leads and shares a positive correlation (43%) with JPMorgan's EMBI index for Ukraine.

Figure 2: Quarterly ASVI for the Search Term "Ukraine IMF"

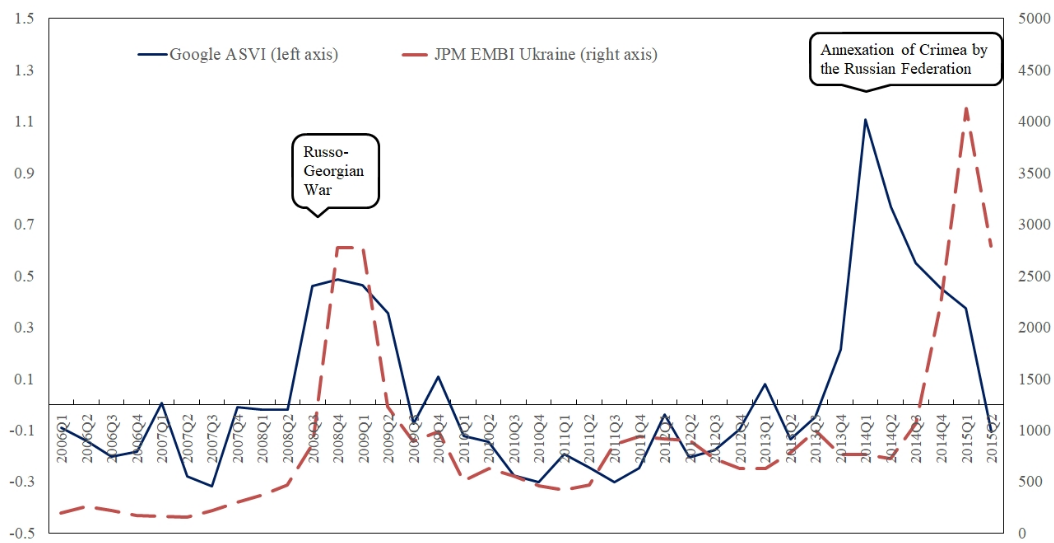

After computing the ASVI, we define an information threshold $\zeta = 0.5 \times \max\{ASVI_t\}$. Our calibration target then becomes the fraction of quarters when Google search intensity about Ukraine is above this threshold. It captures the frequency of large, positive, and relatively discrete jumps in infor-

[18] l_t is not well-identified anyway, since we will be separately calibrating lender wealth. Also even if it is positive, its fixed value would not affect our model mechanism or results. Thus, its specific value is relatively unimportant.

mation acquisition. This is the sort of behavior that our model predicts and so it is a natural target for information cost identification.

For our data, the ASVI threshold $\zeta = 0.55$ and the targeted ratio is 7.1%. It captures the biggest peak in early 2014, which was largely tied to Ukraine's conflict with Russia and the annexation of the Crimean peninsula.

3.1.2. Calibration

To obtain the output process, we estimate via MLE an AR(1) process ($\rho_g = 0.5058, \mu_g = 0.0126, \sigma_g = 0.0846$) on dollar-valued GDP growth for Ukraine from 2004-2014 at a quarterly frequency.

We assume that the risk-free rate is fixed at 1% quarterly; that the lenders exhibit constant relative risk-aversion preferences with CRRA=2, which is standard; and that $\theta = 0.083$, which is an estimate used by Mendoza and Yue (2012) for an average duration of 6 years before returning to international bond market. For simplicity, we also assume that lender wealth is constant over time, i.e., $w_t = w$.

We calibrate the remaining five parameters, $\{\beta, \psi, w, \sigma_m, \kappa\}$, using simulated method of moments (SMM) to target the simulated results from our model at five moments from the corresponding data: Annual default frequency, average debt-service-to-GDP ratio, annual spread volatility, average annual spread, and fraction of time in which information acquisition (IA) is above 50% of its max value.[19]

These parameters are given in Table 1. Each parameter is primarily identified by its corresponding target moment, though there are significant cross-

[19] Our model and its results are at quarterly frequency. The model results are adjusted to annual statistics to match the listed annual targets.

partial effects. The resulting paramaterization is fairly standard[20] and the match is quite close for a 5-dimensional non-linear matching problem.

Table 1: Calibration by Simulated Method of Moments

Parameter	Value	Target Description	Data	Model
Discount factor	β=0.811	Annual Default Frequency	1.5%	1.2%
Output cost (known)	ψ=0.0226	Average Debt-Service-to-GDP Ratio	12.6%	11.3%
Lender wealth	w=2.50	Average Spread	6.5%	5.7%
Unobs shock std dev	σ_m=0.0153	Spread Volatility	5.5%	5.8%
Unit info cost	κ=0.000522	Fraction of Quarters with $IA > \zeta$	7.1%	7.4%

In the following sections, we explore the model's quantitative implications.

3.2. Model Behavior

Before we exposit the properties that are unique to this model, we show that along many dimensions it preserves key features of a standard quantitative sovereign default model. For instance, Figure 3a gives the bond demand functions in equilibrium. It exhibits a simple, downward sloping feature that looks remarkably similar to Aguiar and Gopinath (2006) or Arellano (2008) despite the added complexity of endogenous information acquisition. We can see that better growth shocks lead to higher price schedules.

The equilibrium policy functions are given in Figure 3b. We can see that better growth shocks lead to higher debt issuance. This is a standard feature of the quantitative models discussed in Aguiar et al. (2016a).

Figure 4 provides an event study surrounding a default. Again, the model behaves as a standard model would. We can see that spreads increase prior to

[20]While β may at first glance appear to be very low, it is in the neighborhood of estimates from similar models e.g. Aguiar and Gopinath (2006) or Aguiar et al. (2016a).

Figure 3: Benchmark Behavior: Policy and Pricing Functions

(a) Equilibrium Pricing Function

(b) Equilibrium Policy Function

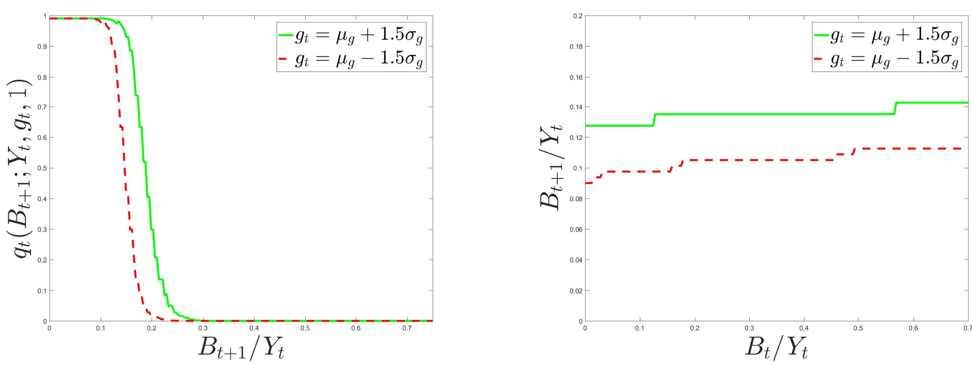

a default. This is because growth follows the boom-bust cycle documented in Aguiar et al. (2016a): Leading up to a default is a series of benevolent growth shocks that induce excessive borrowing, which raises spreads. A default occurs when the sovereign accumulates a large amount of debt in this fashion and then experiences a severe and unanticipated drop in growth (not shown in the figure).

What is novel in this model is the state-contingent acquisition of information: More information is endogenously acquired during times of crises, i.e., near default. Recall from the calibration that the ergodic mean of the debt-to-GDP ratio is about 11.3%. We can see in the policy functions in Figure 5a that for debt levels above this, information precision depends on the underlying growth shock i.e. the observed shock. If growth is high, no useful information is acquired unless debt levels become very large. However, when growth is low, the forecaster acquires useful information and the lenders pay attention to it for lower debt levels as well. Notice further that for very high

Figure 4: Benchmark Behavior Around Default

(a) **(b)**

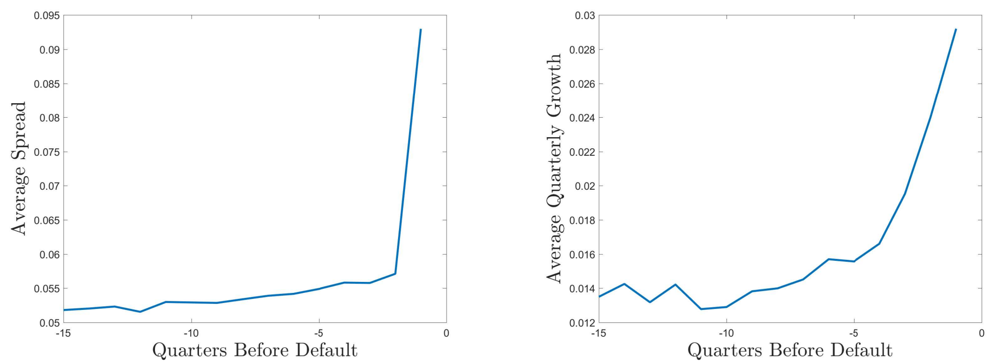

debt levels information acquisition drops off. This is because default is near certain in these regions, regardless of the realization of the unobserved shocks. Consequently, there is no point for the forecaster to pay a cost to learn about those shocks.

The bond demand schedule across x_t can be seen Figure 5b. For low debt levels, the forecaster does not acquire useful information about unobserved shocks; this implies that lenders do not pay attention to these uninformative signals and thus prices do not react. Hence, there is no difference across the two price schedules for debt-to-GDP levels lower than about 10% or so. As debt levels rise, though, the forecaster begins to acquire useful information, the lenders start to pay attention to it, and the bond demand schedules begin to diverge according to the different signals. This divergence lines up exactly with the non-zero information-acquisition regions from Figure 5a.

What do these policy functions imply for the behavior of endogenous objects? Figure 6 provides a graph of optimal signal correlation leading up to a

Figure 5: Information Acquisition: Model Objects

(a) Information Acquisition Policy Function **(b)** Equilibrium Bond Demand Function

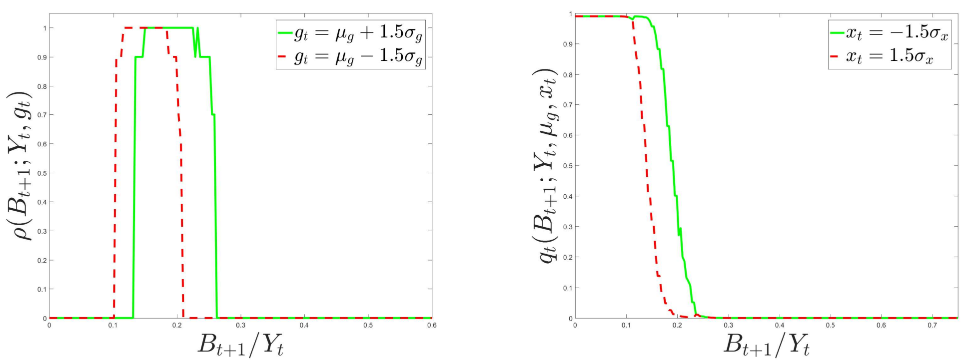

default event. We can see that attention increases before such crises. Signal precision is relatively low in the periods prior to the event, i.e., during non-crisis times. As observed states indicate that default risk is rising, however, the forecaster becomes more aggressive in its acquisition of information, increasing it from an average of about $\rho_{mx} = 0.055$ to $\rho_{mx} = 0.083$ on the cusp of a default event. The acquisition of more information near defaults or crises has a number of interesting implications that we will explore over the next few subsections.

3.3. Comparative Statics

The state-contingent acquisition of information and information cost have large effects on some model moments and minimal effects on others. Among those that are relatively invariant to information costs are average and median debt levels as well as default frequency, all of which hover around their benchmark levels for the complete relevant set of information costs.

Figure 6: Information Acquisition Before Default

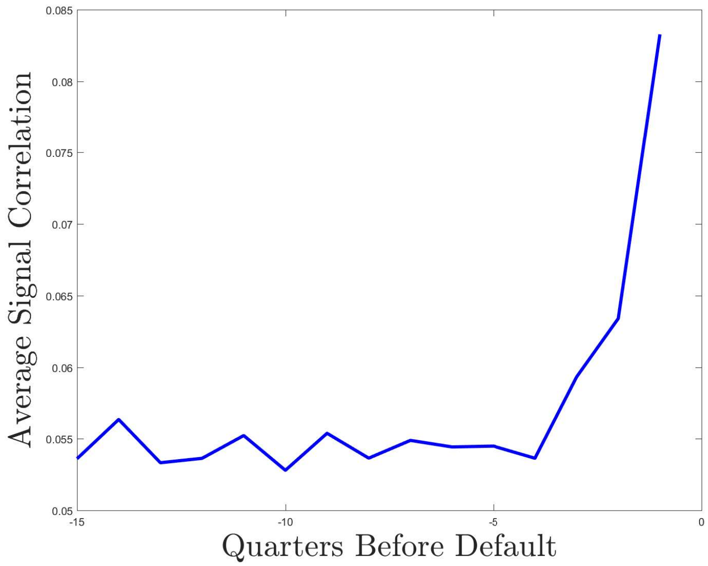

The most relevant affected moments are shown in Figure 7, which provides a comparative static of how the model responds to an array of information costs. These moments consist of the average spread, the median spread, the spread volatility, and the attention fraction. In the figure, black lines correspond to the left y-axis, while blue lines with intermittent crosses correspond to the right y-axis.

As information costs decrease, both the fraction of time spent paying attention to the sovereign and the spread volatility increase significantly. The intuition for the former is trivial and was in fact used for the calibration; the intuition for the latter is that cheaper information means that unobserved shocks are priced more often through the signals instead of being considered at their average. This naturally increases spread volatility.

The impact across measures of central tendency as information costs fall is

Figure 7: Relevant Moments Across κ

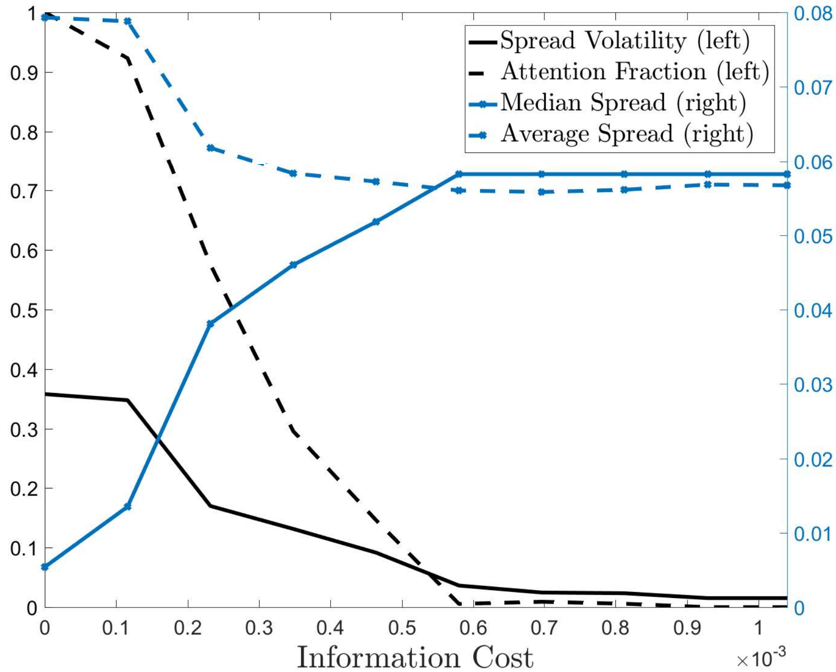

quite interesting. The average spread rises modestly as information costs fall, but the median spread falls substantially. The reason for the former is the increased spread levels during crises. Crises are no more frequent as information costs fall, but during them imminent default risk is more accurately priced so spreads are higher. These equally rare but increasingly large spikes bring the average spread up as information costs fall.

On the other hand, the median spread falls with information costs. This is because the risk premium demanded by the lenders falls as the uncertainty generated by m_{t+1} is more frequently tamed by more accurate information regarding it. And this reduction in risk premium happens in most states of the world i.e. non-crisis times, which brings the median down. This result accords with the findings of Bernoth and Wolff (2008) and Iara and Wolff (2014) that increased transparency reduces borrowing costs.

3.4. Implications and Results

3.4.1. Time-Varying Volatility

The first key result is that our model endogenously generates time-variation in sovereign spread volatility. In the model, spreads exhibit increased volatility during crises when lenders price the unobserved shocks more accurately, rather than considering them to be at their mean, which they do during non-crisis times. Further, the time-varying volatility is strongly countercyclical and positively correlated with spread levels. This implies that one could interpret our framework as a microfoundation for such models with exogenous time-varying volatility as Melino and Turnbull (1990) or Fernández-Villaverde et al. (2011).

To assess our model's ability to generate time-variation in the spread volatility, we propose a model-free metric of the time-variation, which we call the **Crisis Volatility Ratio** or CVR.[21] It is defined as follows: In a series of data, either simulated or empirical, let $\hat{T}$ denote the set of all periods in which the change in the spread from the prior period is above the 97.5 percentile of its distribution. We follow Aguiar et al. (2016a) in calling such events "crises," and by construction they are about 8.3× more likely to happen than default events and can be observed in the data even if a default cannot. With this

[21] Typically, when measuring time-varying volatility, the literature imposes quite a bit of structure. For instance, Melino and Turnbull (1990) and Fernández-Villaverde et al. (2011) can measure the impact of time-varying volatility in the context of a stochastic volatility model. Imposing this or another similar structure to measure the quantitative efficacy of the model would be inappropriate in our case, since we know that the data-generating process for the simulated data is not a stochastic volatility model.

notation, we define the CVR as

$$CVR = \frac{1}{|\hat{T}|} \sum_{t \in \hat{T}} \frac{\hat{\sigma}_{t:t+w}}{\hat{\sigma}_{t-w-1:t-1}}$$

where $\hat{\sigma}_{x:y}$ is the sample standard deviation calculated using the periods from x to y. This ratio compares the volatility in a window of w periods immediately prior to a crisis to the volatility in a window of w periods after. Neither window includes the crisis itself. In the benchmark, we set $w = 5$. If the CVR is larger than one, then crisis periods tend to be more volatile than non-crisis periods. To give a sense of the metric, any AR(1) process would feature an average CVR equal to one since the volatility of the innovation is independent of the any prior observation.

We first compute the CVR for the data (the first column in Table 2). We can see that time-varying volatility is a strong feature of the data by this metric: The CVR in the data is more than 3.5 times what than an AR(1) would suggest. This is no surprise, as strong time-variation in the volatility is documented by Fernández-Villaverde et al. (2011).

Table 2: Crisis Volatility Ratios (First Row) and Counterfactual Moments

Data (Ukraine)	Benchmark Model	$\kappa = \infty$ Counterfactual	$\sigma_m = 0$ Counterfactual
3.67	2.86	1.33	1.27
Def. Freq.	1.2%	1.2%	0.8%
E[Debt Service/GDP]	11.3%	11.3%	13.3%
E[Spread]	5.7%	5.7%	3.7%
Std(Spread)	5.8%	1.5%	2.0%
Fraction with $IA > \zeta$	7.4%	0.0	0.0

What is surprising is the capacity of our model to explain this. Our benchmark model (the second column in Table 2) can explain roughly 78% of this non-targeted moment: 2.86 out of 3.67. This is worth noting, since all time-variation in the volatility is generated endogenously.

The vast bulk of this explanatory power is due to costly information acquisition. To see this, we compare our model to two different counterfactuals and report the CVRs in the last two columns in Table 2. First, we re-solve the model for the same parameters but with infinite information costs, such that unobserved shocks are never priced. Our benchmark model generates a 115% increase in the CVR (from 1.33 to 2.86) relative to the this counterfactual. Second, we re-solve the model with the same parameters but no hidden information i.e. $\sigma_m = 0$. Our model generates a 125% increase in the CVR (from 1.27 to 2.86) relative to this counterfactual, which is equivalent to the model of Aguiar et al. (2016b) with permanent shocks, short-term debt, and a modestly different calibration.

It is worth noting that there is some time-variation in volatility in the model even without information frictions. This is because even those publicly known shocks' fluctuations are not always relevant for default risk and consequently are not always priced. This implies lower price volatility in normal times than crises. This is true in any model of endogenous sovereign default and accords with the findings of Bocola and Dovis (2016). Nevertheless, our model's ability to more than double this underlying the time-variation in spread volatility and thus bring it to empirically relevant levels is substantial and noteworthy.

While our benchmark model is calibrated to match the data and these counterfactuals are not, Table 2 also gives the relevant model moments for these counterfactual models, which are within the empirically relevant neighborhood for a typical emerging market.

Figure 8: Crisis Volatility Ratios Across κ

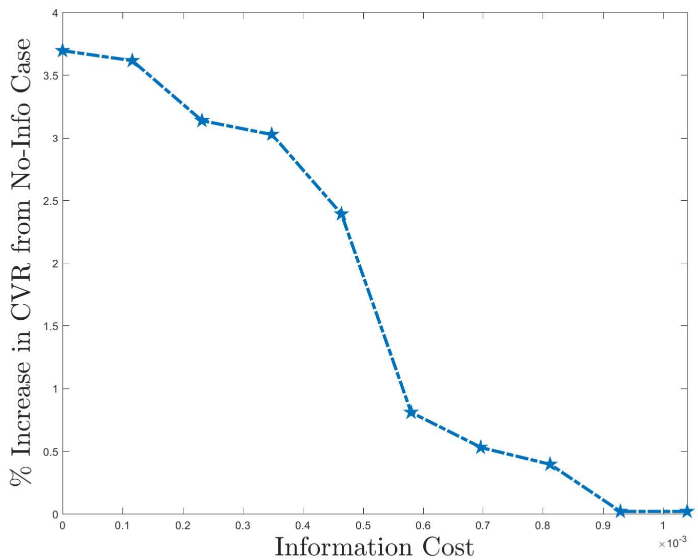

Further evidence of the power of costly information acquisition to generate time-variation in the spread volatility can be found in Figure 8, which gives the percent increase in the CVR from the infinite-cost counterfactual for a wide range of different information costs. It makes clear that the CVR is monotonically decreasing in the information cost. Further, this relationship is quite steep, with the full-information model generating a more than 350% increase over the no-information model.

3.4.2. State-Contingent Risk Premia

The second result we highlight is that the composition of spreads is not the same during crisis times and non-crisis times. To understand this, first note that we can intuitively break the spread on sovereign debt into two categories:

default risk and a risk premium for that default risk.[22]

$$
\begin{aligned}
Sprd_t \;\approx\; & \overbrace{E_t[1 - R_{t+1}(\tilde{s}_{t+1}, \tilde{m}_{t+1}, B_{t+1})]}^{\text{Risk-Neutral Spread/Default Probability}} \;+ \\
& \overbrace{\frac{-cov_t(R_{t+1}(\tilde{s}_{t+1}, \tilde{m}_{t+1}, B_{t+1}), u'(c_L(\tilde{s}_{t+1}, \tilde{m}_{t+1}, B_{t+1})))}{E_t[R_{t+1}(\tilde{s}_{t+1}, \tilde{m}_{t+1}, B_{t+1})]E_t[u'(c_L(\tilde{s}_{t+1}, \tilde{m}_{t+1}, B_{t+1}))]}}^{\text{Risk Premium}}
\end{aligned} \tag{5}
$$

where R is the binary repayment function and $u'(\cdot)$ is the lenders' marginal utility. Notice that the risk premium will always be positive, since $u'(\cdot)$ is a decreasing function and c_L is an increasing function of debt repayment. Thus, the covariance is negative and the overall term is positive.

When lenders pay attention to the forecaster's signal about normally unobserved states, they learn more about the realization of those shocks. In particular, the conditional volatility $\sigma_{m|x}$ shrinks. This pushes down the risk premium term in Equation 5 by reducing the covariance. This occurs even in the case when $x_t = 0$ and the signal merely supports the prior of the lenders, which is that $m_{t+1} = 0$.

This has practical consequences for default risk inference. While default risk is high during a crisis, so too is the investors' attention and their ability to contain that risk since they are acquiring more precise signals about unobserved shocks. This implies that their effective risk-aversion is lower and thus that the risk premium comprises a relatively smaller share of the spread during a crisis than during normal times. Consequently, if an econometrician were not to take this into account, instead employing a more standard sovereign default model with constant information acquisition to infer default risk from

[22] This decomposition follows from a first-order approximation. Derivations can be found in Appendix C.

spread data, she would *underestimate* default risk during crises: She would assume the risk premium to comprise a higher share than it actually does.

Figure 9: Spread Decomposition During Crises

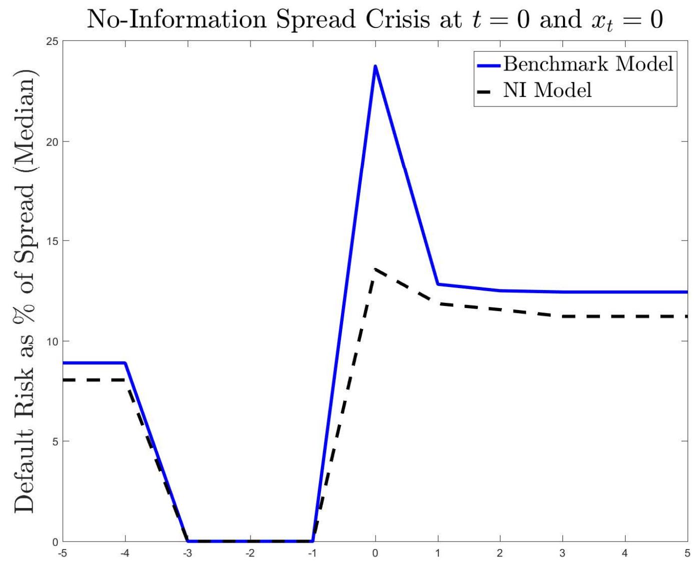

Our model allows us to quantify this compositional shift in the risk spread. To do so, we construct an artificial, non-equilibrium no-information price schedule, which prices the exact same default risk as the benchmark but under the assumption that $\kappa = \infty$. We then compare simulated spreads from the benchmark model to this alternative spread series around a typical crisis.

In particular, we condition on periods that lie in the intersection of two sets: The first is the top 2.5%ile of the spread change distribution in the no-information counterfactual series, which does not price m_{t+1} in any capacity; the second is the set of all periods for which $x_t = 0$. During such events there is no difference in beliefs regarding the average m_{t+1} between the two series, and thus the perceived default risk is the same. However, the perceived volatility of m_{t+1}, and thus the required risk premium, will vary across the two.

We isolate all such crises in a simulation with length 1.5 million periods, and compute the median share of default risk as a fraction of the total spread. The stochastic impulse-response function can be found in Figure 9. During a spread crisis, the benchmark model puts the median default risk share at 23.7%, while the no-information counterfactual puts this figure at 13.6%, a difference of about 10 percentage points. The average difference between the series in non-crisis times, on the other hand, is a mere 0.73 percentage points, an order of magnitude smaller. Thus, default risk *as a share of the total spread* increases by substantially more during crises than a model without costly information acquisition would suggest.

This default-risk composition is particularly interesting since it follows from the fact that risk-premia depend on *country-specific* states. Thus, it cannot be controlled for using global metrics, such as the CBOE VIX or the P/E ratio, as is often done (Aguiar et al. [2016a] or Bocola and Dovis [2016]). Rather, our theory suggests that in order to accurately assess default risk, some metric of forecaster/investor information acquisition, such as SVI or ASVI, must be controlled for.

3.4.3. Transparency

The last result we highlight regards the benefits and costs of transparency in light of our model mechanism. Costly information acquisition can be interpreted in many ways in the context of our model. One way the unit information cost κ can be understood as is the level of transparency that a sovereign has about its domestic affairs and finances. Interpreted this way, our model can also offer some interesting insights for transparency policies.

Figure 10: Sovereign Welfare Across Information Costs

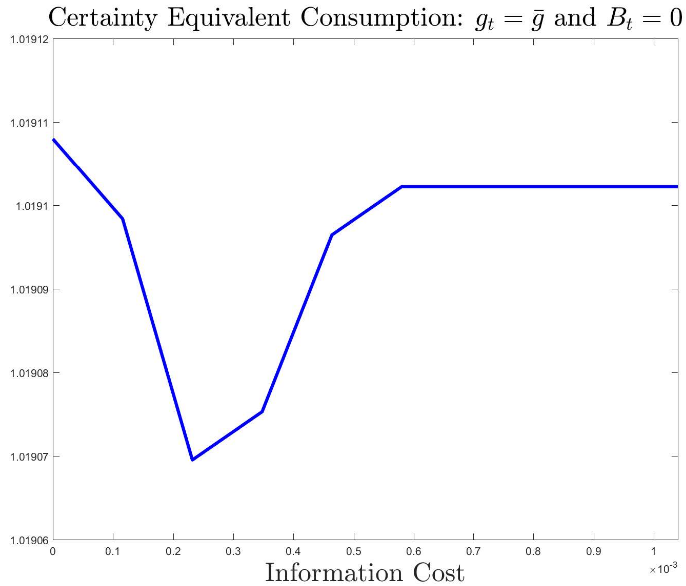

Our model suggests that transparency is a double-edged sword. When there is zero transparency i.e. information is infinitely costly, lenders always demand a risk premium for the unobserved shocks, especially during crisis times. This makes it more expensive for the sovereign to borrow and service debt. Having more transparency will benefit the sovereign by lowering the risk premium. On the other hand, when the sovereign is fully transparent, although risk premium for the unobserved shocks disappears, it is replaced by substantial spread volatility, since now what used to be unobserved shocks are always reflected in bond prices. This will hurt the risk-averse sovereign. This welfare cost of public information precision differs from that of Morris and Shin (2002), which falls instead on investors and results from an overweighting of public information in coordination games.

Therefore, transparency brings about a risk-shifting: There is the benefit of lower risk premia, but also the cost of higher price volatility. To illustrate this insight, we show the sovereign's welfare levels along different information costs in Figure 10, which is evaluated at the steady state growth and at zero debt.[23] We can see that as the information cost decreases from the highest end, the sovereign's welfare first decreases in response to the risk-shifting and higher volatility; but once the cost gets low enough, its welfare eventually reverses course and increases sharply as we approach full information. In this region, the enjoyment afforded by lowered borrowing costs overtakes the excessive volatility and the sovereign is better off.

The model thus suggests that there may be a period of pain associated with a country undertaking greater transparency measures before it reaches the welfare-maximizing level of full information.

4. Conclusion

In this paper, we explored the consequences of costly information acquisition on the pricing of sovereign risk. We constructed and calibrated a structural model of endogenous default and information acquisition, making novel use of Google search data.

We demonstrated that costly information acquisition generates country-specific time-varying volatility in sovereign bond spread; implies a compositional shift in that spread during crises; and highlights both the benefits and costs for emerging markets' welfare with regard to increasing transparency. We also laid the groundwork for information cost identification strategy through relevant attention metrics.

[23]This is the welfare metric used by Chatterjee and Eyigungor (2012). The valley-shape pattern remains the same if we evaluate welfare at the ergodic mean of the debt-to-GDP ratio instead of zero debt.

Possible extensions to our framework could include rollover crises in the vein of Cole and Kehoe (1996), long-maturity debt (Hatchondo and Martinez [2009] or Chatterjee and Eyigungor [2012]), or persistent unobserved shock processes. The intuition of our results would not change with any of these extensions, though the quantitative results may be affected.

References

Aguiar, Mark and Gita Gopinath, "Defaultable Debt, Interest Rates, and the Current Account," *Journal of International Economics*, 2006, *69* (1), 64–83.

___ **, Satyajit Chatterjee, Harold Cole, and Zachary R. Stangebye**, "Quantitative Models of Sovereign Debt Crises," *Handbook of Macroeconomics, Volume II*, 2016.

___ **,** ___ **,** ___ **, and** ___ , "Quantitative Models of Sovereign Debt Crises," *Handbook of Macroeconomics, Volume II (Forthcoming)*, 2016.

Angeletos, George-Marios and Ivan Werning, "Crises and Prices: Information Aggregation, Multiplicity, and Volatility," *American Economic Review*, 2006, *96* (5), 1720–1736.

Arellano, Cristina, "Default Risk and Income Fluctuations in Emerging Economies," *American Economic Review*, 2008, *98* (3), 690–712.

Bacchetta, Philippe and Eric van Wincoop, "Infrequent Portfolio Decisions: A Solution to the Forward Discount Puzzle," *American Economic Review*, 2010, *100* (3), 870–904.

Barber, Brad M. and Terrance Odean, "All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors," *Review of Financial Studies*, 2008, *21* (2), 785–818.

Bassetto, Marco and Carlo Galli, "Is Inflation Default? The Role of Information in Debt Crises," *Mimeo*, 2017.

Bernoth, Kerstin and Guntram B. Wolff, "Fool the Markets? Creative Accounting, Fiscal Transparency, and Sovereign Risk Premia," *Scottish Journal of Political Economy*, 2008, *55* (4), 465–487.

Bi, Huixin and Nora Traum, "Estimating Sovereign Default Risk," *American Economic Review*, 2012, *102* (3), 161–166.

Bloom, Nicholas, "The Impact of Uncertainty Shocks," *Econometrica*, 2009, *77* (3), 623–685.

Bocola, Luigi, "The Pass-Through of Sovereign Risk," *Journal of Political Economy*, 2016, *124* (4), 879–926.

___ **and Alessandro Dovis**, "Self-Fulfilling Debt Crises: A Quantitative Assessment," *Mimeo*, 2016.

Carlson, Mark and Galina B Hale, "Rating Agencies and Sovereign Debt Rollover," *Topics in Macroeconomics*, 2006, *6* (2).

Catao, Luis A.V., Ana Fostel, and Sandeep Kapur, "Persistent Gaps and Default Traps," *Journal of Development Economics*, 2009, *89* (2), 271–284.

Chari, V.V. and Patrick J. Kehoe, "Hot Money," *Journal of Political Economy*, 2003, *111* (6), 1262–1292.

Chatterjee, Satyajit and Burcu Eyigungor, "Maturity, Indebtedness, and Default Risk," *American Economic Review*, 2012, *102* (6), 2674–2699.

Cole, Harold L. and Patrick J. Kehoe, "Models of Sovereign Debt: Partial Versus General Reputations," *International Economic Review*, 1998, *39* (1), 55–70.

___ **and Timothy J. Kehoe**, "A Self-Fulfilling Model of Mexico's 1994-1995 Debt Crisis," *Journal of International Economics*, 1996, *41* (3-4), 309–330.

___ **, Daniel Neuhann, and Guillermo Ordonez**, "Debt Crises: For Whom the Bell Tolls," *NBER Working Paper No. 22330*, 2016.

Da, Zhi, Joseph Engelberg, and Pengjie Gao, "In Search of Attention," *Journal of Finance*, 2011, *66* (5), 1461–1499.

Dow, James and Gary Gorton, "Noise Traders," *NBER Working Paper No. 12256*, 2006.

Durdu, C. Bora, Ricardo Nunes, and Horacio Sapriza, "News and Sovereign Default Risk in Small Open Economies," *Journal of International Economics*, 2013, *91* (1), 1–17.

Eaton, Jonathon and Mark Gersovitz, "Debt with Potential Repudiation: Theoretical and Empirical Analysis," *Review of Economic Studies*, 1981, *48* (2), 289–309.

Fernández-Villaverde, Jesús, Pablo A. Guerron-Quintana, Juan Rubio-Ramirez, and Martin Uribe, "Risk Matters: The Real Effect of Volatility Shocks," *American Economic Review*, 2011, *101* (6), 2530–2561.

Gervais, Simon, Ron Kaniel, and Dan H. Mingelgrin, "The High-Volume Return Premium," *The Journal of Finance*, 2001, *56* (3), 877–919.

Grossman, Sanford, "On the efficiency of competitive stock markets where traders have diverse information," *Journal of Finance*, 1976, *31* (2), 573–585.

___ **and Joseph Stiglitz**, "On the impossibility of informationally efficient markets," *American Economic Review*, 1980, *70* (3), 393–408.

Hatchondo, Juan Carlos, "Asymmetric Information and the Lack of International Portfolio Diversification," *Working Paper University of Rochester*, 2004.

___ **and Leonardo Martinez**, "Long-Duration Bonds and Sovereign Default," *Journal of International Economics*, 2009, *79* (1), 117–125.

Holden, Steinar, Gisle James Natvik, and Adrien Vigier, "Equilibrium Rating and Debt Crises," *International Economic Review*, 2017, *Forthcoming (August).*

Iara, Anna and Guntram B. Wolff, "Rules and Risk in the Euro Area," *European Journal of Political Economy*, 2014, *34*, 222–236.

Johri, Alok, Shahed K. Kahn, and Cesar Sosa-Padilla, "Interest Rate Uncertainty and Sovereign Default Risk," *Mimeo*, 2015.

Justiniano, Alejandor and Giorgio E. Primiceri, "The Time-Varying Volatility of Macroeconomic Fluctuations," *American Economic Review*, 2008, *98* (3), 604–641.

Kopits, George and Jon Craig, "Transparency in Government Operations," *IMF Occasional Paper No. 158*, 1998.

Kyle, Albert S., "Continuous auctions and insider trading," *Econometrica*, 1985, *53* (6), 1315–1335.

Lizarazo, Sandra V., "Default Risk and Risk Averse International Investors," *Journal of International Economics*, 2013, *89*, 317–330.

Lucas, Jr. Robert E., "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 1972, *4* (2), 103–124.

Manso, Gustavo, "Feedback Effects of Credit Ratings," *Journal of Financial Economics*, 2013, *109* (2), 535–548.

Melino, Angelo and Stuart M. Turnbull, "Pricing Foreign Currency Options with Stochastic Volatility," *Journal of Econometrics*, 1990, *45* (1-2), 239–265.

Mendoza, Enrique G. and Vivian C. Yue, "A General Equilibrium Model of Sovereign Default and Business Cycles," *Quarterly Journal of Economics*, 2012, *127* (2), 889–946.

Morris, Stephen and Hyun Song Shin, "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks," *American Economic Review*, 1998, *88* (3), 587–597.

___ **and** ___ , "Social Value of Public Information," *American Economic Review*, 2002, *92* (5), 1521–1534.

Passadore, Juan and Juan Pablo Xandri, "Robust Conditional Predictions in Dynamic Games: An Application to Sovereign Debt," *Job Market Paper*, 2015.

Poterba, James M. and Kim Rueben, "State Fiscal Institutions and the US Municipal Bond Market," *Chapter in NBER Book: Fiscal Institutions and Fiscal Performance*, 1999.

Pouzo, Demian and Ignacio Presno, "Sovereign Default Risk and Uncertainty Premia," *SSRN Working Paper No. 2710469*, 2015.

Sandleris, Guido, "Sovereign Defaults: Information, Investment and Credit," *Journal of International Economics*, 2008, *76* (2), 267–275.

Seasholes, Mark S. and Guojun Wu, "Predictable Behavior, Profits, and Attention," *Journal of Empirical Finance*, 2007, *14* (5), 590–610.

Sedlacek, Petr, "Growth and Uncertainty over the Business Cycle," *CEPR Working Paper No. 11296*, 2016.

Seoane, Hernan D., "Time-Varying Volatility, Default and the Sovereign Risk Premium," *Mimeo*, 2015.

Shannon, Claude E., "Channels with Side Information at the Transmitter," *IBM Journal of Research and Development*, 1958, *2* (4), 289–293.

Sims, Christopher A., "The Implications of Rational Inattention," *Journal of Monetary Economics*, 2003, *50* (3), 665–690.

——, "Rational Inattention: Beyond the Linear-Quadratic Case," *American Economic Review*, 2006, *96* (2), 158–163.

Sosa-Padilla, Cesar, "Sovereign Defaults and Banking Crises," *MPRA Paper No. 41074*, 2012.

Stangebye, Zachary R., "Dynamic Panics: Theory and Application to the Eurozone," *MPRA Working Paper No. 69967*, 2015.

Van Nieuwerburgh, Stijn and Laura Veldkamp, "Information Immobility and the Home Bias Puzzle," *Journal of Finance*, 2009, *64* (3), 1187–1215.

Veldkamp, Laura, *Information Choice in Macroeconomics and Finance*, Princeton University Press, 2011.

Appendix A. Proof of Isomorphism Between Preference and Cost Shocks

Notice that under log-preferences, we can express the value of default as

$$\begin{aligned} V_{D,t}(s_t, m_t) = & \log\left(\hat{Y}_t\right) \; + \; m_t \; + \\ & \beta E_{\tilde{s}_{t+1}|s_t} \left[\phi V_t(\tilde{s}_{t+1}, 0, 1) + (1-\phi) V_{D,t}(\tilde{s}_{t+1}, 1)\right] \end{aligned}$$

where $\hat{Y}_t = Y_t e^{-\psi}$ is output net of default costs. Now we can define $\hat{V}_{D,t}(s_t) = V_{D,t}(s_t, m_t) - m_t$, where $\hat{V}_{D,t}(s_t)$ is the value of default when the only cost of defaulting is $\phi(g_t)$.

With this notation, the default decision can be written as

$$d_t(m_t, s_t, B_t) = \mathbf{1}\{V_{R,t}(s_t, B_t) < \hat{V}_{D,t}(s_t) + m_t\}$$

In this isomorphic environment, m_t becomes a preference shock to the utility of defaulting.

Appendix B. Solving the Model

We consider the case of CRRA preferences for generality. We stationarize this model by dividing the sovereign resource constraint in every period by Y_t, and the value functions by $Y_t^{1-\gamma}$. This delivers a convenient recursive structure which is independent of both time and the level of output. We denote $b' = B_{t+1}/Y_t$ and $c = C_t/Y_t$.

$$\begin{aligned} v_R(g, b) = & \max_{b' \geq 0} \; E_{\tilde{m}', \tilde{x}} \left[\frac{c(\tilde{x})^{1-\gamma}}{1-\gamma} + \beta E_{\tilde{g}'|g} e^{\tilde{g}'(1-\gamma)} v(\tilde{g}, b', \tilde{m}')\right] \\ & \text{s.t.} \;\; c(\tilde{x}) = 1 - b e^{-g} + q(g, b', \tilde{x}) b' \end{aligned}$$

The value of default is scaled similarly, yielding

$$\begin{aligned} v_D(g, m) = & \frac{e^{(-\psi+m)(1-\gamma)}}{1-\gamma} + \\ & \beta E_{\tilde{g}', \tilde{m}'|g} \left[\phi e^{\tilde{g}'(1-\gamma)} v(\tilde{g}', 0, 0) + (1-\phi) e^{\tilde{g}'(1-\gamma)} v_D(\tilde{g}', 0)\right] \end{aligned}$$

This stationarization implies that we can express the default policy function using only stationarized model objects, since Y_t does not influences the default decision once g_t is known.

$$d(m, g, b) = \mathbf{1}\{e^{-g(1-\gamma)} v_R(g, b) < e^{-g(1-\gamma)} v_D(g, m)\}$$

The benchmark model with sovereign's log-preferences will simply be the limiting case as $\gamma \to 1$. This will imply that the stationarized model will feature log flow utility and that the impact of g on the effective discount factor vanishes.

The forecaster's problem is already stationarized, since it deals only with the distributions. We can stationarize the lenders' problem as well under the assumption that $w_t = wY_t$.

$$\max_{b'_i} \; E_{\tilde{m}', \tilde{g}' | x_i, g} \left[\frac{c_i'^{1-\gamma_L}}{1-\gamma_L} \right]$$
$$\text{s.t.} \;\; c'_i = (w - b'_i q)(1+r) + b'_i \left[1 - d(\tilde{m}', \tilde{g}', b')\right]$$

We solve the model using a discretized grid over the state space: We approximate the growth process using a Tauchenized Markov process with 25 grid points; we use a uniform grid over debt levels from [0, .75] of 201 points; and we discretize m_t, x_t, and $\rho_{mx,t+1}$ across 11 points each. Relevant model moments remain virtually unchanged when we increase the size of this grid along any dimension.

Appendix C. Spread Decomposition

For simplicity of notation, we drop the time subscripts and work with the recursive notation. We first consider the risk-neutral case. Here, the lenders' first-order condition together with the market clearing condition implies the following pricing expression:

$$q(B'|s, x) = \frac{E_{\tilde{s}', \tilde{m} | s, x}[R(\tilde{s}', \tilde{m}, B')]}{1+r}$$

where $R(\tilde{s}', \tilde{m}, B')$ is the binary repayment function of the sovereign. The effective interest rate is $\hat{r}_{RN} = 1/q$. Taking a log of the previous expression yields:

$$\begin{aligned}
&\log(1 + \hat{r}_{RN}) = \log(1 + r) - \log(E_{\tilde{s}', \tilde{m}|s,x}[R(\tilde{s}', \tilde{m}, B')]) \\
\implies &\hat{r}_{RN} - r \approx -E_{\tilde{s}', \tilde{m}|s,x}[\log(R(\tilde{s}', \tilde{m}, B'))] \\
\implies &sprd_{RN} \approx E_{\tilde{s}', \tilde{m}|s,x}[1 - R(\tilde{s}', \tilde{m}, B')]
\end{aligned}$$

The second line follows from the expectation of a first-order approximation around the mean. The third follows from the first-order approximation $x \approx \log(1 + x)$ applied to $R(s, m, B) = 1 - D(s, m, B)$, assuming default risk to be relatively small.

To compute the overall spread, we use a similar strategy. Note the first-order necessary condition is

$$q(B'|s, x) = \frac{E_{\tilde{s}', m|s,x}[R(\tilde{s}', \tilde{m}, B')u'(c_L(\tilde{s}', \tilde{m}, B'))]}{(1 + r)E_{\tilde{s}', m|s,x}[u'(c_L(\tilde{s}', \tilde{m}, B'))]}$$

i.e., the ratio of the expected marginal utilities in repayment states over the expected marginal utilities in all states. Following a similar procedure as before, we arrive at

$$\begin{aligned}
\log(1 + \hat{r}) - \log(1 + r) = &- \log(E_{\tilde{s}', \tilde{m}|s,x}[R(\tilde{s}', \tilde{m}, B')]) - \\
&\log\left(1 + \frac{cov(R(\tilde{s}, \tilde{m}, B'), u'(c_L(\tilde{s}, \tilde{m}, B')))}{E[R(\tilde{s}, \tilde{m}, B')]E[u'(c_L(\tilde{s}, \tilde{m}, B'))]}\right)
\end{aligned}$$

If both the default risk and the covariance expression are relatively small (close to zero), then a first-order approximation yields Equation 5.

Appendix D. Robustness: Alternative Measure of Attention and Alternative Search Terms

In this section we explore the robustness of our proxy for information acquisition, Google SVI.

Appendix D.1. Alternate Attention Metrics

First, we explore the correlation of our chosen information acquisition metric, SVI, with another common metric of investor attention in the literature, extreme daily returns (Barber and Odean [2008]). In particular, we use daily return data on the stripped sovereign spread from JP Morgan's EMBI database and compute the largest monthly return in absolute value. We then compare this series to the monthly SVI series in Figure D.1. While they do not line up perfectly, there is much co-movement ($\rho = 0.3541$) and thus our metric lines up reasonably well with this other proxy.

Figure D.1: Comparison of SVI and Extreme Returns

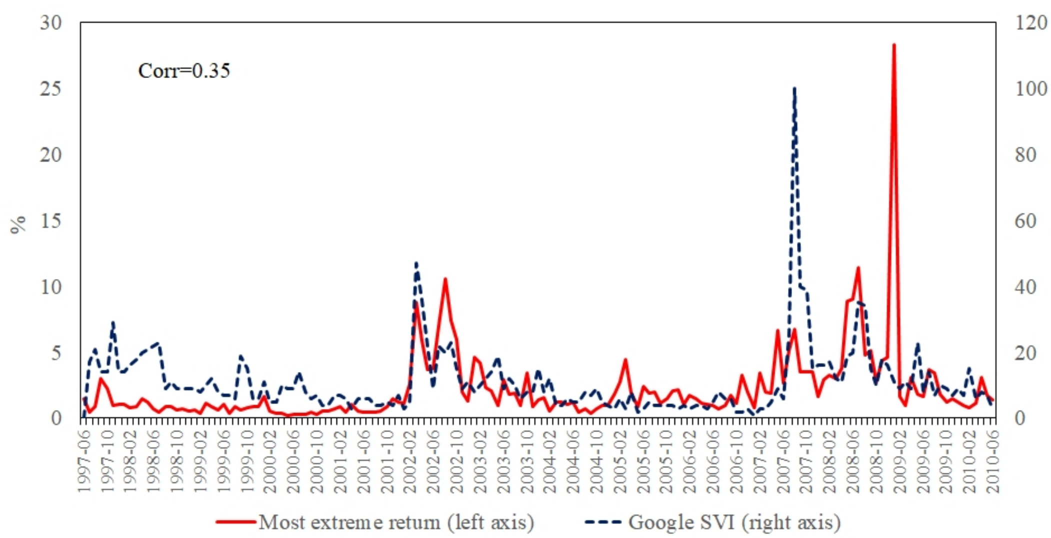

Most of this co-movement is driven by extreme negative returns, which is consistent with our theory since investors pay attention more during crises.

The correlation with SVI and the minimum daily return in a month is -0.5208, which provides even further evidence of the relevance of our metric.

Another alternative considered in the literature are daily trade volumes. However, such data is difficult to acquire and may not exist for sovereign bonds. This is the case for our benchmark choice of Ukraine, for which almost all trades are executed over-the-counter instead of in an exchange.

Figure D.2: Comparison of SVI and Rating Changes

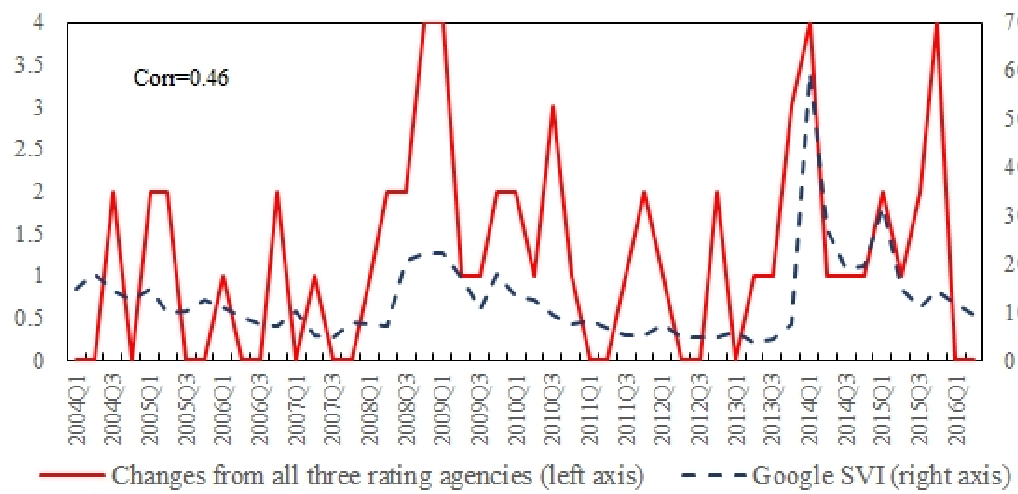

Finally, given that credit rating agencies may serve the role of forecaster, we also check the correlation between our SVI metric and the frequency of credit rating changes for a measure of forecaster information acquisition (Figure D.2). We collect the credit rating from the three major credit rating agencies (S&P, Moody's, and Fitch) and sum the number of rating changes in each quarter. We then compare this series to our SVI series. The correlation is 0.46.

Appendix D.2. Alternate Search Terms/Languages

Since it is not entirely clear what the appropriate search term should be, we consider a handful of alternate search terms to ensure that we are capturing as closely as possible investor attention and not merely general inquiry searches. In addition to the benchmark term, 'Ukraine IMF,' we also consider the terms 'Ukraine bloomberg' and 'Ukraine Reuters.' The series are juxtaposed in Figure D.3. We can see from the figure that the benchmark term has an extra 'spike' during the Russo-Georgian war that the alternative terms do not. This is not a problem for us, however, since this attention spike was not large enough to place it over the mid-point, and thus our identification strategy did not count this as a period in which attention was paid. Rather, for all three terms, the only relevant spike in the sample occurred around the Russian annexation of Crimea in 2014. Further, our target calibration statistic does not vary greatly for these terms, with the fraction of high attention for 'Ukraine bloomberg' being 4.3% and for 'Ukraine reuters' being 4.4%. While these come in lower than our benchmark target of 7.1%, Figure 7 (our comparative static exercise) reveals that the quantitative impact this cost difference would have on our simulated moments and quantitative results is small.

Figure D.3: Comparison of Benchmark Search Term to Alternate Search Terms

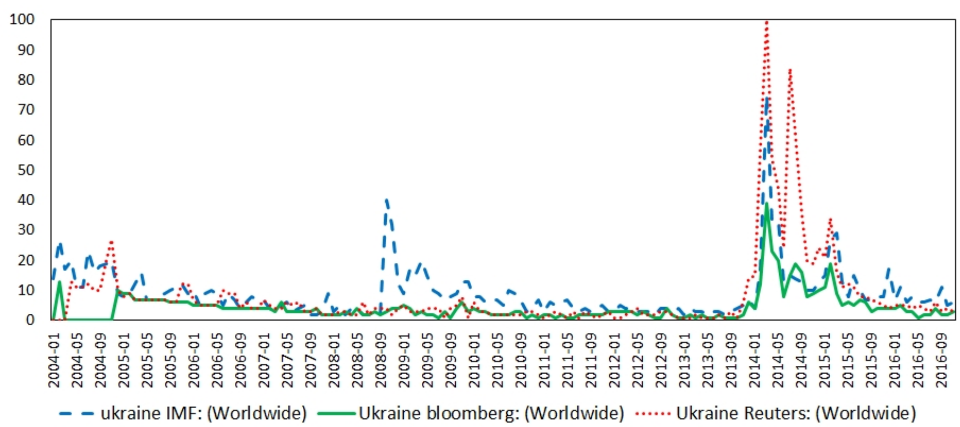

We also consider our benchmark search in other languages. Figure D.4 provides a color-coded world map showing the language of maximum search volume for each country for the three most common languages: English, Russian, and Chinese. We can see that English is far and away the most dominant language for these searches, even dominating search volume over Russian in Russia. The only exception is in China, where Chinese is dominant.

Figure D.4: Benchmark Search Language versus Most Common Alternatives

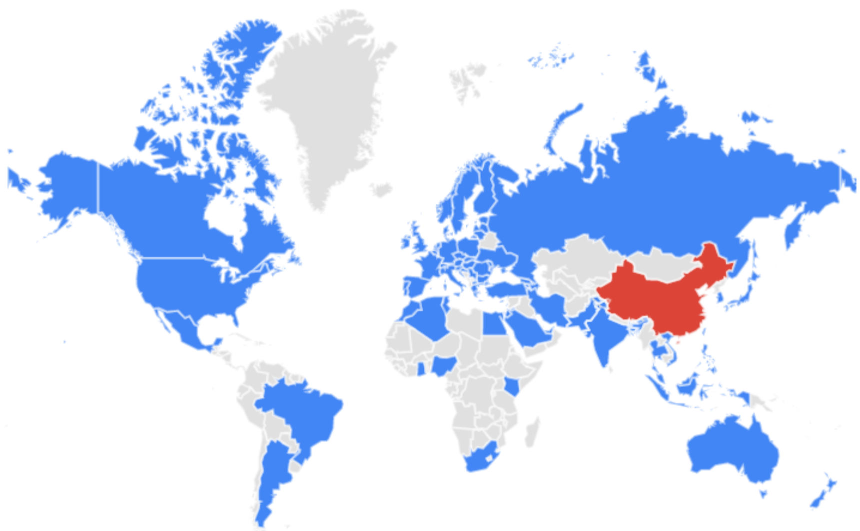

Blue: English (Benchmark), Yellow: Russian, Red: Chinese